



De-homogenization using Convolutional Neural Networks

Elingaard, Martin O.; Aage, Niels; Bærentzen, J. Andreas; Sigmund, Ole

Published in:
Computer Methods in Applied Mechanics and Engineering

Link to article, DOI:
[10.1016/j.cma.2021.114197](https://doi.org/10.1016/j.cma.2021.114197)

Publication date:
2022

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Elingaard, M. O., Aage, N., Bærentzen, J. A., & Sigmund, O. (2022). De-homogenization using Convolutional Neural Networks. *Computer Methods in Applied Mechanics and Engineering*, 388, Article 114197. <https://doi.org/10.1016/j.cma.2021.114197>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



De-homogenization using convolutional neural networks

Martin Ohrt Elingaard^{a,b}, Niels Aage^a, Jakob Andreas Bærentzen^b, Ole Sigmund^{b,*}

^a Department of Mechanical Engineering, Technical University of Denmark, Kgs. Lyngby, Denmark

^b Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kgs. Lyngby, Denmark

Received 28 April 2021; received in revised form 14 September 2021; accepted 14 September 2021

Available online xxxx

Abstract

This paper presents a deep learning-based de-homogenization method for structural compliance minimization. By using a convolutional neural network to parameterize the mapping from a set of lamination parameters on a coarse mesh to a one-scale design on a fine mesh, we avoid solving the least square problems associated with traditional de-homogenization approaches and save time correspondingly. To train the neural network, a two-step custom loss function has been developed which ensures a periodic output field that follows the local lamination orientations. A key feature of the proposed method is that the training is carried out without any use of or reference to the underlying structural optimization problem, which renders the proposed method robust and insensitive wrt domain size, boundary conditions, and loading. A post-processing procedure utilizing a distance transform on the output field skeleton is used to project the desired lamination widths onto the output field while ensuring a predefined minimum length-scale and volume fraction. To demonstrate that the deep learning approach has excellent generalization properties, numerical examples are shown for several different load and boundary conditions. For an appropriate choice of parameters, the de-homogenized designs perform within 7–25% of the homogenization-based solution at a fraction of the computational cost. With several options for further improvements, the scheme may provide the basis for future interactive high-resolution topology optimization.

© 2021 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Keywords: Deep learning; Convolutional neural networks; Topology optimization; De-homogenization

1. Introduction

In many ways the fields of computer vision and topology optimization are closely related, e.g. in both cases, the optimization domain constitutes a set of elements/pixels in a grid, and the problem can be formulated as minimizing a functional to obtain an optimized configuration. In computer vision, an optimized configuration could be a set of segmentation labels, which in the context of topology optimization [1] could be thought of as material phases. Due to the closeness of the tasks, it is only reasonable to assume that some of the impact which deep learning has had on the field of computer vision will transfer to the field of topology optimization.

Led by this promise, recent years have seen a surge in publications on the application of deep learning [2] to topology optimization within a variety of problems such as minimum compliance [3–9], thermal compliance [10–12], micro-structure design [13–17] and generative design [18]. Despite these efforts, deep learning has yet to

* Corresponding author.

E-mail address: sigmund@mek.dtu.dk (O. Sigmund).

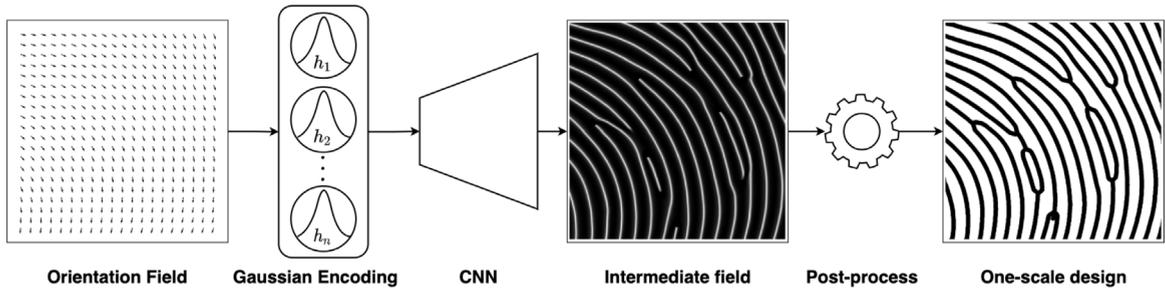


Fig. 1. Data pipeline from orientation field to intermediate density field on to one-scale design.

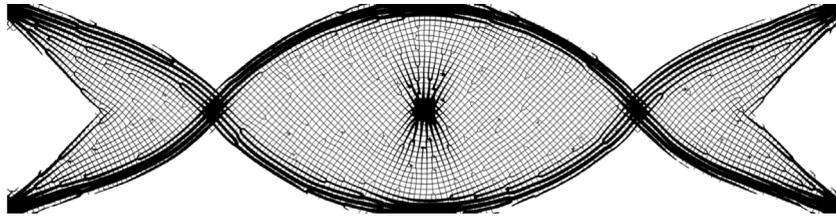


Fig. 2. De-homogenization of a 200×50 homogenization-based solution onto a fine mesh of 8000×2000 elements. Compliance of the de-homogenized design is within 17% of the homogenization solution, adheres to a predefined minimum relative thickness, and has been de-homogenized in just 15 s on a modern day laptop.

make a major impact within topology optimization, and only performs at a similar level to traditional topology optimization methods at very low design resolutions, or for highly restricted problem domains and boundary conditions, where the cost of generating a dataset and training a neural network is not prohibitive. Put bluntly, and due to the above reasons, none of the so far published deep learning end-to-end methods allow for a level of generalization which makes it advisable to use them for practical applications. Moreover, for end-to-end machine learning to have a significant impact on the field of topology optimization, it is clearly a requirement that it is not only faster and scales better, but also that the solutions produced are of comparable quality to those computed using more conventional optimization methods [1]. Furthermore, this must be true even when the network is provided inputs that are very different from those in the training data. This appears to be a very challenging problem, but attractive alternatives to end-to-end learning present themselves. For example, it is generally much easier to incorporate constraints and ensure a physically sound design in the classical mathematical programming methods, where only one design problem is considered at a time, whereas convolutional neural networks have shown state-of-the-art performance on tasks such as image denoising [19] and upsampling [20] where local information plays an important role. Hence, a combined method, in which the two approaches complement each other, seems to be a much more viable direction to pursue [21].

Here we report on a deep learning approach that allows for problem definition generalization and which is capable of producing high resolution and detailed designs. The key to this achievement lies in the fact that the proposed method does not provide an end-to-end topology optimization approach. Instead, it provides an alternative to traditional de-homogenization approaches [22–25] by using the orientations from a homogenization-based topology optimization solution [23,26] as inputs for a fully convolutional neural network [27], which upsamples them to an intermediate field $\tilde{\rho}$. The network is trained using a synthetic dataset generated by sampling the spatial gradient orientations from a field of low-frequency sines, and an encoding based on splitting the angular information into multiple channels is used to ease learning for the network. A custom loss function ensures that the intermediate field is periodic according to a predefined wave-length and that its spatial gradients are orthogonal to the input directions. Furthermore, two regularization terms are included that favor a solution with sufficient spatial variation, and where branching mainly takes place in the solid phase. Fig. 1 shows the data pipeline from the orientation field to the intermediate field and subsequent post-processing. The post-processing consists of a series of fast steps, all running in linear time, which project the desired lamination widths on the intermediate field to yield the de-homogenized

one-scale design. An example of the outcome of the suggest learning-based de-homogenization approach applied to a double clamped cantilever case is shown in Fig. 2.

We remark, that solving the homogenization-based topology optimization problem is already very fast, i.e. the problem is close to convex and much information can be extracted, even from coarse finite element meshes [25]. The current bottleneck in de-homogenization methods is the expensive solving of one large-scale least square problem per lamination direction. In this work, we propose a learning-based solution to this problem that promises significant time savings. With computation times around 2 s for projections onto a 2400×1200 mesh on a standard laptop and lots of potential for further improvements and porting to GPUs, this approach may pave the way for future real-time de-homogenization in the iterative homogenization-based optimization loop [28].

The structure of the paper is as follows; in Section 2 the relevant methods are outlined including homogenization-based topology optimization, network architecture, loss function, and post-processing scheme. Section 3 presents the synthetic dataset generation and input encoding. In Section 4 results for several different minimum compliance examples are presented and compared to a reference solution provided by the homogenization-based method. Finally, a discussion of the strengths and weaknesses of the proposed method and potential improvements are discussed in Section 5.

2. Methods

This section provides a brief introduction to homogenization-based topology optimization, followed by a detailed description of the proposed neural network and lamination width projection scheme.

2.1. Homogenization-based topology optimization

The seminal papers by Bendsøe and Kikuchi [26,29] introduced the so-called homogenization-based topology optimization approach. Crucially, this method overcame the ill-posedness of element or node-based design parameterizations by introducing optimal, multi-scale, so-called rank-N (or slightly sub-optimal rectangular hole) microstructures, oriented along principal stress directions. This approach was considered quite complex and also resulted in “gray” solutions that could not be manufactured. Therefore, and specifically after the simpler density-approach was proven physically admissible [30], homogenization-based topology optimization was largely abandoned in favor of simpler approaches. Now, where simpler density approaches have reached giga-scale resolution [31] that require access to supercomputers, and where advances in additive manufacturing techniques allow realization of efficient multi-scale structures, there has been a renewed interest in homogenization-based topology optimization approaches, that can be run efficiently on simple computers and who’s output can be realized through novel de-homogenization approaches [22–25].

The homogenization-based approach to topology optimization can be explained through the visualization in Fig. 3. A design domain (a) is discretized into a relatively low number of elements and each element has three design variables. Lamination parameters μ_1 and μ_2 denote fractions of solid material at two different length scales and θ determines the orientation. For this so-called rank-2 material, homogenized effective properties can be determined in closed form (see e.g. [1]). Based on analytical gradient computations, these three design fields can be optimized, resulting in e.g. the cantilever beam structure in Fig. 3a, where gray color indicates optimized element-wise volume fractions and blue and red lines indicate direction and proportions of the two local lamination thicknesses. The exact formulation of the homogenization-based topology optimization approach used here follows next.

2.1.1. Optimization problem

The goal of the considered homogenization-based topology optimization problem is to minimize the structural compliance \mathcal{C} , i.e. the work done by external forces, subject to a volume constraint. This objective is augmented by two additional terms; an angle regularization term \mathcal{F}_θ [25] and a penalization of the solid area [32]. The design variables are the two element-wise lamination parameters μ_1 , μ_2 , the angle θ , as well as an auxiliary material indicator field s .

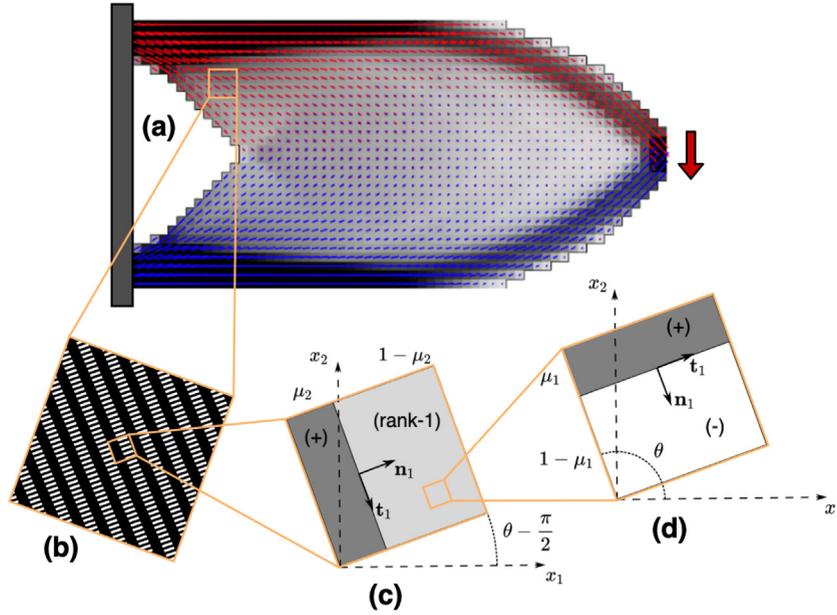


Fig. 3. Parameterization of rank-2 microstructures in homogenization-based topology optimization. (a) Cantilever beam with boundary conditions. (b) Rank-2 laminate structure. (c) Parameterization of rank-2 microstructure. (d) Parameterization of rank-1 microstructure made from a stiff material (+) and a weak material (-). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The full optimization problem, including constraints, is stated as

$$\begin{aligned}
 \min_{\mu_1, \mu_2, \theta, \mathbf{s}} & : \frac{C(\mu_1, \mu_2, \theta, \mathbf{u}) + \gamma_\theta \mathcal{F}_\theta(\theta)}{C_{ref}} + \Gamma \frac{V_{\Omega, \text{struct}}^e(\mathbf{s})}{V_\Omega} \\
 \text{s.t.} & : \mathbf{K}(\mu_1, \mu_2, \theta) \mathbf{u} = \mathbf{f} \\
 & : \mathbf{v}^T \boldsymbol{\rho}^d(\mu_1, \mu_2) - V_{max} \leq 0 \\
 & : \mu_{min} \leq \mu_1, \mu_2 \leq 1 \\
 & : -4\pi \leq \theta \leq 4\pi
 \end{aligned} \tag{1}$$

Here \mathbf{K} is the stiffness matrix, \mathbf{f} is the force vector, \mathbf{u} is the solution to the equilibrium equations, \mathbf{v} is a vector of element volumes, $\boldsymbol{\rho}^d$ is a vector of element densities, V_{max} is the maximum amount of allowed material and C_{ref} is a reference compliance. Finally, γ_θ and Γ are constants used to scale the importance of each term in the objective and μ_{min} is the minimum bound on the admissible layer thickness, from here on referred to as the minimum relative thickness. Box constraints on the angle variables θ are made large enough never to become active by setting them to $\pm 4\pi$.

The relation between design variables and physical variables is

$$\tilde{\boldsymbol{\mu}}_i^k = (\alpha_{out} + (1 - \alpha_{out}) \tilde{\boldsymbol{\mu}}_i) \tilde{\mathbf{s}}^k \tag{2}$$

$$\boldsymbol{\rho}^k = \tilde{\boldsymbol{\mu}}_1^k + \tilde{\boldsymbol{\mu}}_2^k - \tilde{\boldsymbol{\mu}}_1^k \tilde{\boldsymbol{\mu}}_2^k \tag{3}$$

$$V_{\Omega, \text{struct}}^e(\mathbf{s}) = \int_{\Omega} \tilde{\mathbf{s}}^e d\Omega, \tag{4}$$

where $\alpha_{out} = 10^{-9}$ is a small value indicating the stiffness in the void phase, and Ω denotes the optimization domain. Here, $\tilde{\square}$ denotes filtered variables and $\tilde{\square}$ denote projected variables. The subscript $k \in \{e, b, d\}$ refers to the eroded e , blue print b and dilated b designs of the robust formulation [33]. With this formulation, we ensure clear designs without non-physical lamination parameters in the interval $[0, \mu_{min}[$. For a detailed description of this approach, the readers are referred to the paper by Giele et al. [32].

2.2. Convolutional neural network

The homogenization-based topology optimization solution contains a set of lamination parameters describing the underlying microstructure in each element, but does not in itself provide a physically realizable solution, as infinite periodicity of the microstructure is assumed within each element (Fig. 3b). Together with the element-wise varying lamination angles, this poses the difficult challenge of seamlessly connecting the microstructure in adjacent elements to form a mechanically sound and efficient de-homogenized design at finite periodicity. In the context of previous de-homogenization approaches, this is where the current bottleneck lies [22–24]. The most expensive part of existing de-homogenization approaches is the solution of a least-squares problem on an intermediate mesh to identify two auxiliary fields with gradients corresponding to the normals of the two lamination directions. Hence, it is interesting to investigate if this task can be performed more efficiently with a learning-based approach.

The most straightforward way to utilize a neural network for solving this task would be to do it in a supervised manner [34]. In the supervised learning case, the network takes the lamination parameters as input and outputs a de-homogenized design which is then evaluated against a ground-truth de-homogenized design generated using existing de-homogenization methods [23,24], i.e. the neural network G creates a mapping

$$\boldsymbol{\rho} = G(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \boldsymbol{\theta}) \quad (5)$$

and is evaluated using an error estimate, such as the mean-squared error, between the predicted and ground-truth design

$$\mathcal{L}_{MSE} = \frac{1}{n} \sum_{i=1}^n (\boldsymbol{\rho} - \boldsymbol{\rho}_{gt})^2 \quad (6)$$

where $\boldsymbol{\rho}_{gt}$ denotes the ground-truth design.

In practice, there are several downsides to such an approach. First of all, a very large amount of computation is required in order to generate a dataset of ground-truth designs with sufficient resolution and quality. Secondly, even though de-homogenization methods have improved tremendously in recent years most methods still require some degree of manual intervention, which is both cumbersome and undesirable to do for thousands of samples. Also, an entirely target-based loss function only ensures pixel-wise accuracy, but does not ensure the structural integrity of the design as a whole. Thus, if a few crucial pixels are predicted with low accuracy the design may contain structural disconnects, but exhibit a low geometrical error overall. Of these two issues, the dataset generation is the most pressing, as minor structural disconnects can often be fixed by a subsequent post-processing scheme. Therefore, and due to the large amount of computation and manual work needed to make a supervised approach work, an unsupervised approach [35] is taken in this paper. Our solution is based on the observation that the mechanical soundness of the structure is provided by the homogenization-based solution. Thus, our convolutional network needs only to produce structures of a given periodicity from the lamination orientations provided by the homogenization-based solution, whereas the lamination widths are projected onto the design in a post-processing step. This problem greatly resembles texture synthesis which crucially allows us to train the network using synthetic fields. However, initial tests using the orientation field directly as input for the neural network did not show promising results, and thus an alternative encoding based on splitting the angular field into multiple channels each modeled by a Gaussian was developed as detailed in Section 3.

Given the activations in each channel the network G performs the following mapping

$$\tilde{\boldsymbol{\rho}} = G(\mathbf{H}) \quad (7)$$

where \mathbf{H} is the Gaussian activations, and $\tilde{\boldsymbol{\rho}}$ is an intermediate field whose gradients follow lamination orientations, i.e. $|\nabla \tilde{\boldsymbol{\rho}} \cdot \tilde{\mathbf{e}}_{\theta}| \approx 0$. Note that for a rank-2 microstructure description the lamination orientations correspond to the principal stress orientations.

A fully convolutional neural network [27] is chosen as the desired network architecture as it allows an arbitrary input size, and requires much fewer learning parameters compared to a fully connected neural network. In the bottom layer of the network, a series of Residual Neural Network (ResNet) blocks [36] are used to extract local features from the input orientations. To accurately represent the orientation information in the intermediate field a total upsampling factor of $\times 8$ is performed. Each upsampling layer consists of a nearest-neighbor upsampling followed by two convolutional layers, which have proven less prone to checkerboard-artifacts compared to other

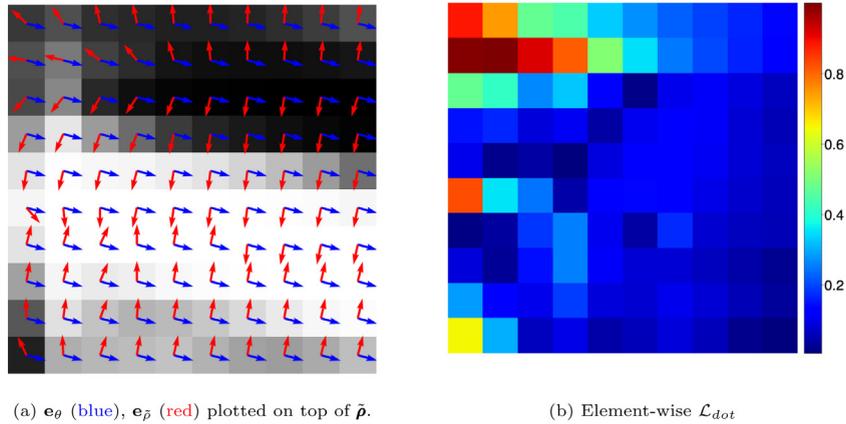


Fig. 4. Visual representation of the dot-product loss near the edge of the domain when the network has been trained. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

upsampling methods such as pixel-shuffle or transposed convolution [37]. The final output layer consists of a 3×3 kernel followed by a Sigmoid activation to scale the output between 0 and 1. More information on the exact network architecture can be found in [Appendix B](#).

To ensure that the output of the neural network follows the lamination orientations a loss function based on the dot-product between the orientations from the homogenization-based solution and the normalized gradients of the intermediate field is used

$$\mathcal{L}_{dot} = \hat{\mathbf{e}}_{\tilde{\rho}}(\mathbf{x}) \cdot \bar{\mathbf{e}}_{\theta}(\mathbf{x}) \quad (8)$$

where \mathbf{e}_{θ} is the vectorized version of the lamination angle, and $\hat{\mathbf{e}}_{\tilde{\rho}} = \nabla \tilde{\rho} / \|\nabla \tilde{\rho}\|_2$ is calculated using the Sobel operator. To avoid very small gradient magnitudes from influencing numerical stability during training a mean filter is applied to $\nabla \tilde{\rho}$ before normalization. [Fig. 4](#) shows a visual representation of the dot-product loss for a small patch of $\tilde{\rho}$ near the edge of the domain. It can be seen that generally, the loss is higher near the edge as the Sobel operator does not provide an accurate estimate of the gradient here, while orientations further from the domain boundaries are near orthogonal. To counteract this issue, replication padding is added to the neural network input.

The dot product loss function \mathcal{L}_{dot} alone provides no control of frequency or periodicity of the output. For practical reasons, a certain length-scale (wavelength) of the microstructure is preferred. In the original work on de-homogenization [23], this was obtained by introducing an average length-scale parameter ε . However, due to divergence of the stress field, this may result in largely varying wavelengths throughout the design domain. To ensure a more well-defined wavelength, also avoiding very long and slender features (and thus indirectly ensuring better local buckling resistance), [38] introduced splitting of bars through forking. Similarly, to achieve an output field with a specific periodicity here, we introduce an additional loss term that favors a certain local wavelength.

To control the periodicity of the output field a windowed Fourier transformation

$$\tilde{f}(x, y) = \mathcal{F}(\tilde{\rho}(x, y) \cdot w(x, y)) \quad (9)$$

is utilized as the basis for the additional term in the loss function. Here a Hamming window w has been used to account for spectral leakage [39]. This additional term measures the energy inside a specified frequency band compared to the total energy of the spectrum

$$E_{\omega} = \frac{\sum_{i_{\omega}} \sum_{j_{\omega}} |\tilde{f}_{i_{\omega}, j_{\omega}}|^2}{\sum_i \sum_j |\tilde{f}_{i, j}|^2} \quad (10)$$

where i_{ω}, j_{ω} index frequencies in the desired range.

In terms of de-homogenization it is most intuitive to think of this frequency band as describing the desired wave-length ε . Since an exact fulfillment of the desired wave-length leads to a very narrow frequency band, which

is impractical during optimization, it is desirable to assign a small fixed width of b pixels to the frequency band. Given a square image of width H the frequency interval is simply

$$\{i_\omega, j_\omega\} = \left\{ (i, j) \mid \frac{H}{\varepsilon} - \frac{b}{2} < \sqrt{i^2 + j^2} < \frac{H}{\varepsilon} + \frac{b}{2} \right\} \quad (11)$$

where we assume the zero-frequency is at the origin. The spectral loss is included in the total loss function as

$$\mathcal{L}(\mathbf{e}_{\tilde{\rho}}, \mathbf{e}_\theta, \tilde{\mathbf{f}}) = \mathcal{L}_{dot}(\mathbf{e}_{\tilde{\rho}}, \mathbf{e}_\theta) - \lambda_\omega E_\omega(\tilde{\mathbf{f}}) \quad (12)$$

Here λ_ω is a weight factor and a minus is used in front of E_ω as we seek to maximize the energy inside the specified frequency band, but we seek to minimize the total objective.

If no regularization is used, the objective in Eq. (12) tends to favor solutions with very low spatial variation in terms of magnitude. This might lead to numerical instabilities when performing normalization of the image gradient orientations, especially since the computations are performed on a single-precision GPU. Two different approaches have been tested in order to counteract an overly smooth representation of $\tilde{\rho}$. The simplest approach is to add weak Gaussian noise to the input orientation fields, before performing the input encoding. This helps the network learn more robust representations of the intermediate field, essentially pushing the values towards the extremes of the Sigmoid function, where small variations have less impact. In practice the Gaussian noise alone is often enough to ensure sufficient spatial variation. In addition to the Gaussian noise a total variation term

$$\mathcal{V}(y) = \sum_{i,j} (y_{i+1,j} - y_{i,j})^2 + (y_{i,j+1} - y_{i,j})^2 \quad (13)$$

is added to the loss to enforce a predefined total variation of the intermediate field.

It is important to note that the total variation term on its own would not provide an explicit way of controlling the periodicity of the intermediate field, as the total variation may be increased in multiple ways, i.e. either through increasing the contrast, the periodicity or a mixture of both. Thus, the total variation and spectral losses should be used in conjunction

$$\mathcal{L}(\tilde{\rho}, \mathbf{e}_{\tilde{\rho}}, \mathbf{e}_\theta, \tilde{\mathbf{f}}) = \mathcal{L}_{dot}(\mathbf{e}_{\tilde{\rho}}, \mathbf{e}_\theta) - \lambda_\omega E_\omega(\tilde{\mathbf{f}}) + \lambda_\tau \left\| \frac{\mathcal{V}(\tilde{\rho})}{\tau} - 1 \right\|_2 \quad (14)$$

where τ is the desired total variation, and λ_τ is a weight factor.

Due to divergence in the stress fields, the local wavelength must be increased or branches must occur. This is unrelated to the underlying homogenization-based topology optimization formulation. In earlier works, we either allowed divergence as also done in the original Michell theory [23], or we introduced branching to prevent long, thin and buckling prone features to appear [38]. Since the periodicity is fixed by the spectral loss, branching is the only option in the present case. Without any regularization, areas around a branching point often take intermediate values. This is undesirable if a thresholding approach is later used to obtain a fully solid-void design as it might lead to disconnects. Thus, it is preferable to drive branching points towards the solid phase. One possible way to do this is to use the dot-product loss in Eq. (8) as an indicator function for branching. This is based on the observation that peaks in the dot-product loss most often coincide with branching, see Fig. 5.

Using the dot-product loss surface directly as an indicator function introduces an unnecessary amount of noise to the optimization, as there are lots of small peaks in the loss-surface not related to branching. Furthermore, the error will often be large near the edge of the domain, as the Sobel operator provides a poor estimate of the gradients in this area. To create a cleaner indicator function, the loss surface is first smoothed using a small Gaussian kernel g

$$\mathbf{I}_{dot}(x, y) = \mathcal{L}_{dot}(x, y) * g(x, y) \quad (15)$$

Subsequently, all values in \mathbf{I}_{dot} closer than 5 pixels to the edge of the domain are set to zero. Now, a loss term favoring branching in the solid phase may be formulated based on the indicator function

$$\mathcal{L}_b(\mathbf{I}_{dot}, \tilde{\rho}) = \mathbf{I}_{dot} \odot (1 - \tilde{\rho}) \quad (16)$$

where \odot indicates element-wise multiplication, and $1 - \tilde{\rho}$ is used such that only dot-product errors in the void-phase are penalized. This naturally drives the intermediate field towards values very close to one, but still far enough apart to provide meaningful image gradients.

Including the branching loss in Eq. (14) leads to the final loss function used for this work, i.e.

$$\mathcal{L}(\tilde{\rho}, \mathbf{e}_{\tilde{\rho}}, \mathbf{e}_{\theta}, \tilde{\mathbf{f}}, \mathbf{I}_{dot}) = \mathcal{L}_{dot}(\mathbf{e}_{\tilde{\rho}}, \mathbf{e}_{\theta}) - \lambda_{\omega} E_{\omega}(\tilde{\mathbf{f}}) + \lambda_{\tau} \left\| \frac{\mathcal{V}(\tilde{\rho})}{\tau} - 1 \right\|_2 + \lambda_b \mathcal{L}_b(\mathbf{I}_{dot}, \tilde{\rho}) \quad (17)$$

where λ_{ω} , λ_{τ} and λ_b are scaling parameters, used to weigh the importance of the spectral, total variation and branching loss in the full objective.

However, minimizing all four terms in Eq. (17) simultaneously often leads to poor local minima, and requires careful balancing of the four objectives to obtain useful results. Based on experience gained from extensive numerical experiments, we remedy this issue by performing the training in two steps: in the first step the network is trained with only the orientation, spectral and total variation losses enabled, i.e. $\lambda_b = 0$. This ensures that the generated intermediate field follows the input orientations, and the prescribed periodicity. In the second step the spectral loss is disabled, and the branching loss is enabled instead, i.e. $\lambda_{\omega} = 0$. The network is then initialized with the weights from the first training step and trained until convergence.

This procedure makes sense as the branching loss relies on branches being the main source of error in the orientation loss, which is not the case during early stages of the first training phase. On the other hand the spectral loss disturbs the generation of well-defined branches, as they violate the prescribed periodicity. The overall training procedure is outlined in Algorithm 1.

Algorithm 1: Two-step training procedure

```

1 First step - orientation, spectral and total variation losses
2  $\mathbf{w}_1 \leftarrow \text{KaimingInit}()$ 
3  $\mathcal{L}_1(\tilde{\rho}, \mathbf{e}_{\tilde{\rho}}, \mathbf{e}_{\theta}, \tilde{\mathbf{f}}) = \mathcal{L}_{dot}(\mathbf{e}_{\tilde{\rho}}, \mathbf{e}_{\theta}) - \lambda_{\omega} E_{\omega}(\tilde{\mathbf{f}}) + \lambda_{\tau} \left\| \frac{\mathcal{V}(\tilde{\rho})}{\tau} - 1 \right\|_2$ 
4 for  $i \leftarrow 1$  to  $n_{1,epochs}$  do
5    $\mathbf{H}, \mathbf{e}_{\theta} \leftarrow \text{DataLoader}(i)$ 
6    $\tilde{\rho} \leftarrow G(\mathbf{H})$ 
7    $\mathbf{e}_{\tilde{\rho}} \leftarrow \text{ImgGrad}(\tilde{\rho})$ 
8    $\tilde{\mathbf{f}} \leftarrow \text{DFT}(\tilde{\rho})$ 
9    $\mathbf{w}_1^{(i)} \leftarrow \mathbf{w}_1^{(i-1)} - \eta \nabla \mathcal{L}_1(\tilde{\rho}, \mathbf{e}_{\tilde{\rho}}, \mathbf{e}_{\theta}, \tilde{\mathbf{f}})$ 
10 end
11 Second step - orientation, total variation and branching losses
12  $\mathbf{w}_2 \leftarrow \mathbf{w}_1$ 
13  $\mathcal{L}_2(\tilde{\rho}, \mathbf{e}_{\tilde{\rho}}, \mathbf{e}_{\theta}, \mathbf{I}_{dot}) = \mathcal{L}_{dot}(\mathbf{e}_{\tilde{\rho}}, \mathbf{e}_{\theta}) + \lambda_{\tau} \left\| \frac{\mathcal{V}(\tilde{\rho})}{\tau} - 1 \right\|_2 + \lambda_b \mathcal{L}_b(\mathbf{I}_{dot}, \tilde{\rho})$ 
14 for  $i \leftarrow 1$  to  $n_{2,epochs}$  do
15    $\mathbf{H}, \mathbf{e}_{\theta} \leftarrow \text{DataLoader}(i)$ 
16    $\tilde{\rho} \leftarrow G(\mathbf{H})$ 
17    $\mathbf{e}_{\tilde{\rho}} \leftarrow \text{ImgGrad}(\tilde{\rho})$ 
18    $\mathbf{I}_{dot} \leftarrow \text{GaussKernel}(\tilde{\rho})$ 
19    $\mathbf{w}_2^{(i)} \leftarrow \mathbf{w}_2^{(i-1)} - \eta \nabla \mathcal{L}_2(\tilde{\rho}, \mathbf{e}_{\tilde{\rho}}, \mathbf{e}_{\theta}, \mathbf{I}_{dot})$ 
20 end
```

2.3. Lamination width projection

The homogenization-based topology optimization is based on optimally oriented rank-2 microstructures as described in Section 2.1. However, to achieve practically realizable single-scale microstructures, we convert the rank-2 microstructures to a single scale rectangular-hole microstructures with little loss in stiffness [40] as discussed in [23].

Given the intermediate field $\tilde{\rho}_1$ for a single lamination direction a post-processing procedure is applied in order to project the desired lamination width onto the intermediate field. Below, the post-processing procedure for a single direction is outlined, but the procedure remains entirely the same for the other direction.

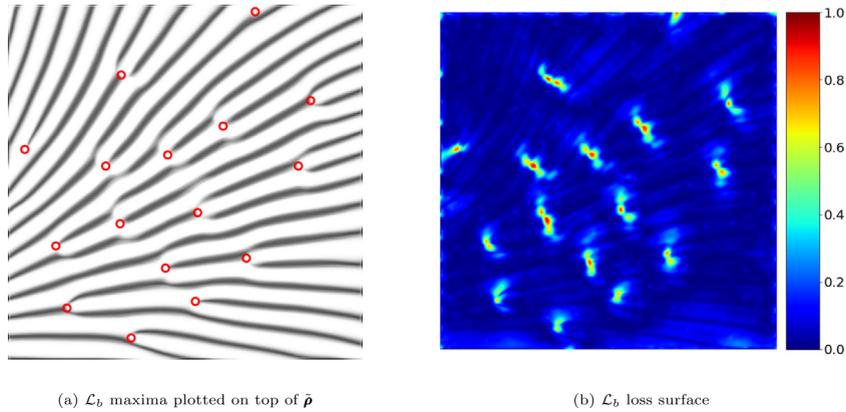


Fig. 5. An intermediate field $\tilde{\rho}$ and the corresponding branching loss surface \mathcal{L}_b . Red circles indicate maxima of the \mathcal{L}_b loss surface.

To make sure that the smallest features in $\tilde{\rho}_1$ are resolved by at least three pixels a bilinear upsampling from the intermediate mesh \mathcal{T}^i to a finer mesh \mathcal{T}^f is first performed. An estimate of the required upsampling factor needed to ensure that \mathcal{T}^f is able to resolve the minimum feature size can be calculated as

$$m_{up} = \left\lceil \frac{h_{min}}{\varepsilon_i \cdot \mu_{min}} \right\rceil \tag{18}$$

where μ_{min} is the minimum relative thickness, h_{min} is the minimum feature size in pixels and ε_i is the wave-length in pixels/period on the intermediate mesh \mathcal{T}^i .

Once $\tilde{\rho}_1$ is of reasonable resolution the skeleton (c.f. Fig. 6c) of the solid phase is extracted. However, in order to do so a few preparation steps are made. First, $\tilde{\rho}_1$ is normalized to the range 0–1 as the network output does not always utilize the entire range of the Sigmoid. Second, branching regions are solidified by extracting all maxima above a certain threshold in the dot-product loss surface, and setting all pixels in a radius of $\varepsilon_i/4$ to 1, c.f. Fig. 6b. Lastly, the upsampled intermediate field is converted to binary image by setting all values above the mean to 1. Once the preparations steps have been performed, a skeletonization [41] is performed on the binary image to obtain the underlying structure of the intermediate field. Since skeletonization thins the solid part of the density field to one-pixel thickness, a dilation with a pixel-radius of 1 is applied to avoid point connections in diagonal bars. The dilated skeleton now serves as the minimum thickness for all members in the structure.

Having established the underlying skeleton and minimum thickness of all members in the structure a width is now assigned to each member. Here a distance transform D_1 (c.f. Fig. 6d) is used to turn the skeleton image into a distance field. Given the distance field a thickness can be assigned to each member of the skeleton by using the lamination width μ_1 to adaptively threshold the distance field

$$\rho_1(\mathbf{x}) = \mathbb{H}(\mu_1(\mathbf{x}) - D_1(\mathbf{x})) \tag{19}$$

where $\mathbb{H}(\mathbf{x})$ is the Heaviside step function.

There are a few caveats to this procedure. First off, due to branching the periodicity in $\tilde{\rho}_1$ is not entirely uniform. This means that if D_1 is normalized to the 0–1 range outliers will influence the normalization, and cause an erroneous projection. To remedy this issue a histogram on D_1 can be used to identify outlier values, and determine an appropriate clipping value before normalizing the distance field. Secondly, μ_1 should be clipped according to $\frac{h_{min}}{P \cdot m_{up}}$ and normalized subsequently, such that the minimum feature size corresponds to the minimum width resolvable by the mesh. Fig. 6 gives an overview of the entire post-processing procedure from the intermediate field to the 1-directional density field.

Once the density fields for both lamination orientations have been calculated the global density field can be obtained as the union between the two fields

$$\rho(\mathbf{x}) = \min \{ \rho_1(\mathbf{x}) + \rho_2(\mathbf{x}), 1 \} \tag{20}$$

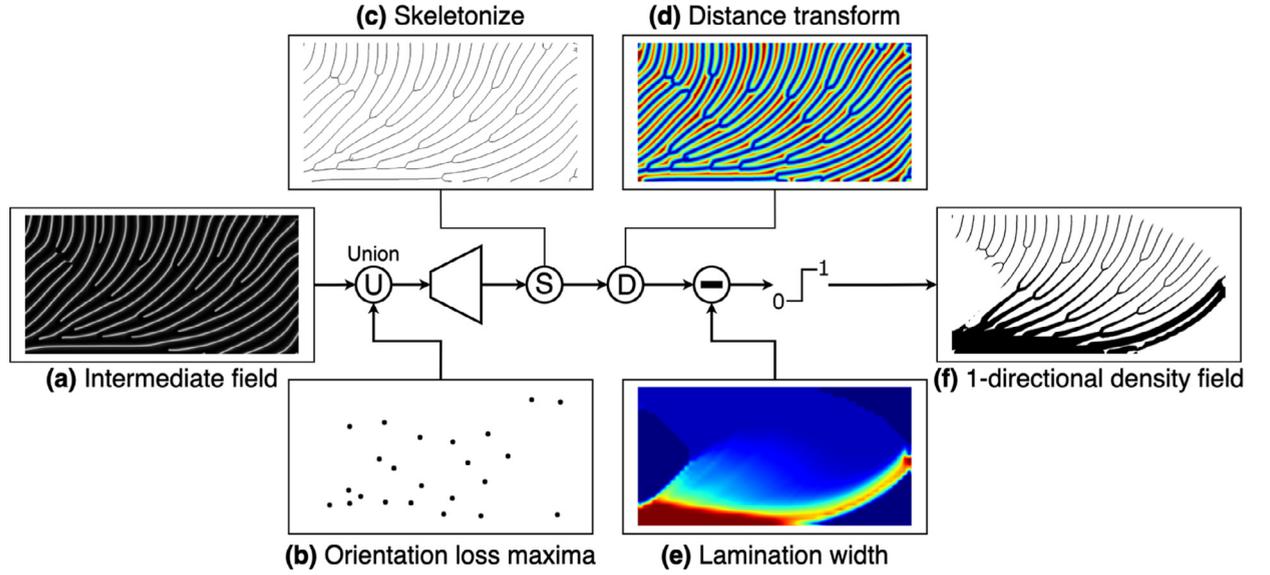


Fig. 6. Graphical illustration of the post-processing procedure.

3. Dataset and input encoding

One of the absolute strengths of the method proposed in this paper, is that training data can be generated synthetically. When training a neural network, using data from the desired application domain, in most cases, leads to the best performance. In our case this type of data would correspond to a set of orientation fields generated using the homogenization-based approach. However, such orientation fields are both expensive to obtain, and might contain singularities, especially around regions where load and boundary conditions have been applied. These singularities often result in non-integrable vector fields, which are not directly suited for de-homogenization [42]. To avoid singularities and lessen the computational burden a synthetic dataset is used instead.

To obtain a synthetic orientation field, we first create a synthetic scalar field, F , by summing products of low frequency sine functions and obtain the needed vector field by taking the gradient:

$$F(x, y) = \sum_n \sum_m c_{n,m} \sin\left(\frac{n\pi x}{x_L}\right) \sin\left(\frac{m\pi y}{y_L}\right) \quad (21)$$

$$v_x(x, y) = \frac{\partial F(x, y)}{\partial x} \quad (22)$$

$$v_y(x, y) = \frac{\partial F(x, y)}{\partial y} \quad (23)$$

where $c_{n,m}$ are individual wave amplitudes. Since the gradient magnitude is of no importance to the orientation, the gradients are normalized to yield

$$\mathbf{e}(\mathbf{x}) = \mathbf{v}(\mathbf{x}) / \|\mathbf{v}(\mathbf{x})\|_2 \quad (24)$$

From the global orientation field generated using Eqs. (21)–(24) smaller patches with different aspect ratios may be sampled. To ensure a somewhat smooth orientation field within each patch a constraint on the maximum angular change θ_{max} between neighboring orientations is introduced. This constraint is enforced by resampling a patch if the constraint is violated, and also ensures that the singularities around local extrema in the global field are avoided. Furthermore, the global orientation field is resampled every N_{rs} iterations with a random number of sine contributions from the set $C = \{6, 8, 10\}$. Fig. 7 shows an example of four patches sampled using the above mentioned approach, while Table 1 provides an overview of the parameters used for the sampling procedure.

From Fig. 7(b) it can be seen that the variation within each patch is generally not that large. Thus, to reduce computation and memory consumption the orientation vectors are subsampled to half the resolution using a

Table 1
Parameters used to sample the synthetic dataset.

Global field size	Patch sizes	C	θ_{max}	N_{rs}
H=W=800	80×80	$n = m = 6$	25 deg	100
	60×120	$n = m = 8$		
	40×160	$n = m = 10$		

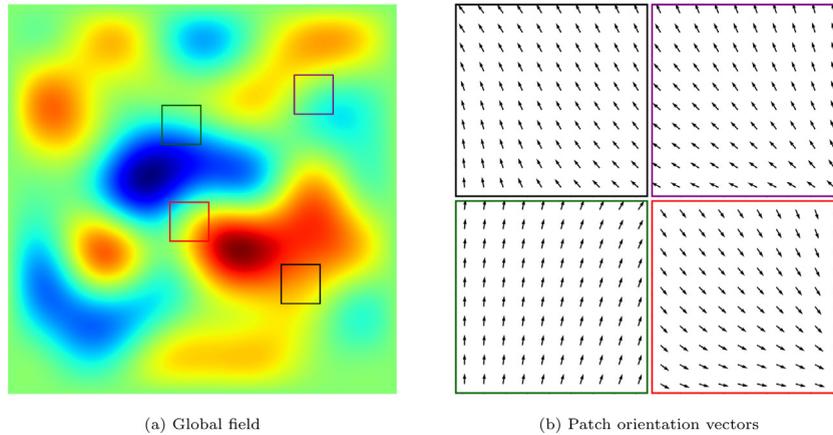


Fig. 7. The global field composed of low-frequency sines and four randomly sampled orientation vector patches. Notice that orientation vectors have been subsampled to improve readability of the plot.

2×2 block average. Furthermore, the block average makes the approach robust towards zero-divisions during the normalization in Eq. (24). Notice, even though the orientation vectors are subsampled it remains important to use a relatively large patch size to capture sufficient variation in the global field.

Initial tests using the orientation field \mathbf{e} directly as input for the neural network did not show promising results, and thus an alternative encoding based on splitting the angular input into N different channels was used. Each channel is now modeled using a Gaussian

$$h = \exp\left(-\frac{(x - c)^2}{r^2}\right) \quad (25)$$

with centers c equally spaced in the interval $[0; \pi]$ and a kernel radius of $r = 2\pi/N$. For this particular case x indicates the input angle.

The encoding constitutes a mapping $\mathbb{R}^1 \rightarrow \mathbb{R}^N$ from the angular information of \mathbf{e} to an activation in each of the N channels and is applied in an element-wise manner. This type of input encoding bears resemblance to binning, but ensures a smooth representation of the activation in each channel as the contribution of the angular input in each channel is weighted by the Gaussian. Fig. 8 shows the activations in each of the channels for $N = 12$ and an orientation field with a close to 90° phase shift.

There are a few caveats which must be taken into account when performing this type of encoding. First, \mathbf{e} is a 2-directional field, and thus invariant to rotations of π . This means that any angular values outside the $[0; \pi]$ range, in which the Gaussians are defined, can simply be shifted to the equivalent value inside the range. Second, to preserve the π -periodicity any activation in one end of the interval must be properly reflected in the other end of the interval. This is done by introducing a set of support Gaussians which extends the $[0; \pi]$ interval by $3r$ in each end. Any activations in these support functions are mapped to the appropriate channel in the opposite end of the interval.

4. Numerical examples

The proposed method generalizes to a wide range of minimum compliance problems as the lamination fields and angles provided by the homogenization-based topology optimization method implicitly take boundary and load

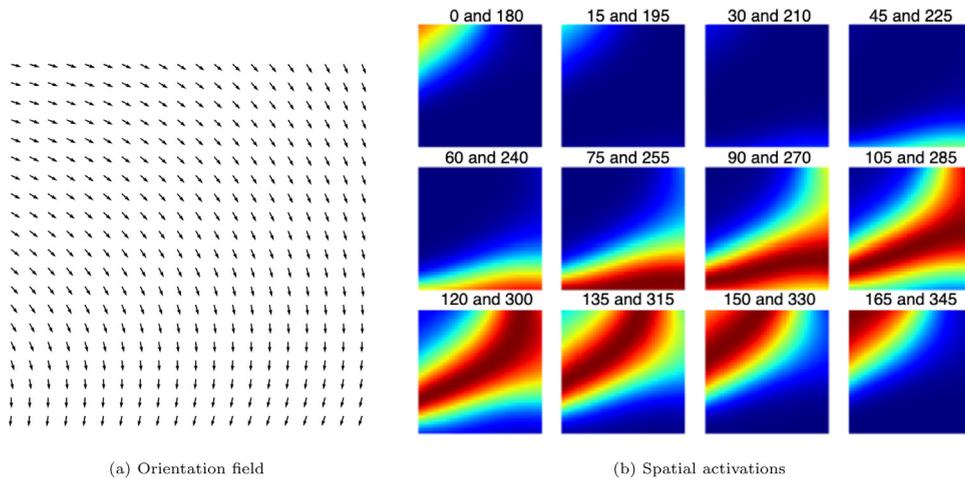


Fig. 8. An orientation field and the corresponding spatial activations in each of the input encoding channels.

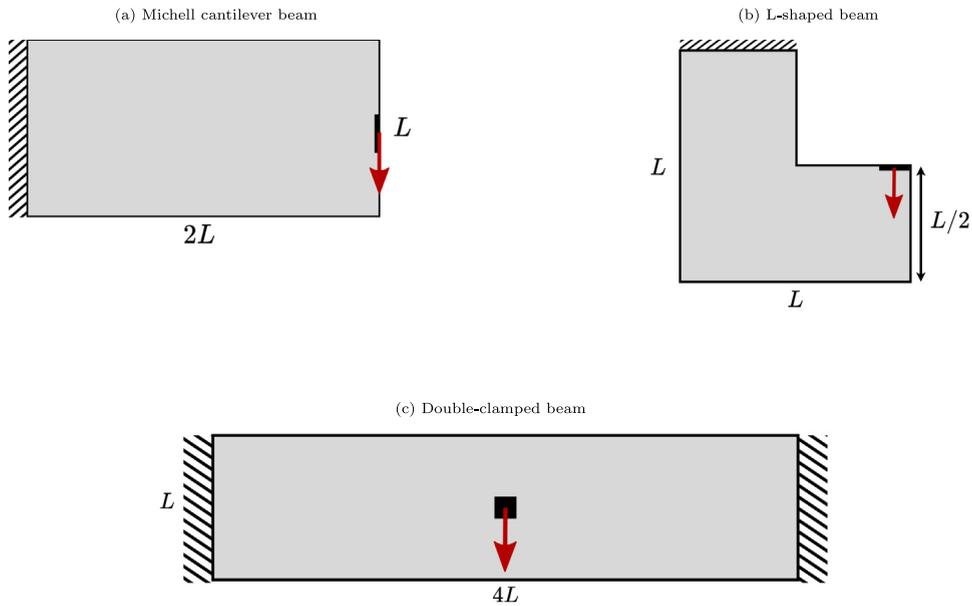


Fig. 9. Design domain and loading for the Michell cantilever, L-shaped and double-clamped beam.

conditions into account. To demonstrate this capability the method is tested on three different problems; the Michell cantilever beam, the double-clamped beam and the L-shaped beam. See Fig. 9 for domain sizes and exact placement of boundaries and loads for the three cases. To avoid stress singularities at load locations the force is applied along a line of $\frac{1}{10}L$ in the Michell cantilever case, a line of $\frac{1}{20}L$ in the L-shaped beam case and on a block of $\frac{1}{10}L \times \frac{1}{10}L$ in the double-clamped beam case. Furthermore, all elements with a density of $\rho > 0.99$ and in a small radius around the load are set to solid. To generate the lamination widths μ_1, μ_2 used during the post-processing scheme and the angles θ serving as input for the neural network the optimization problem stated in Eq. (1) is solved until convergence with $\gamma_\theta = 0$, $\Gamma = 0.05$, and a filter radius of $r_{min} = 1.2$ for varying mesh sizes, volume fractions and relative minimum thicknesses μ_{min} . For all of the numerical examples presented below, replication padding with a width of 2 pixels has been added to the angular input for the neural network, to provide a better estimate of the image gradient near domain boundaries.

Table 2

Performance and computational cost of neural network based de-homogenization approach on the Michell cantilever beam for a 60×30 input. Here h_c is the element size on the coarse mesh (homogenization mesh), ε_i is the wave-length on the intermediate mesh (output of neural network), μ_{min} is the minimum relative thickness, V_{ref} is the reference volume fraction from the homogenization solution, C_{ref} is the reference compliance from the homogenization solution, h_f is the element size on the fine mesh (de-homogenization mesh), ε_f is the wave-length on the fine mesh, V_f is the volume fraction of the de-homogenized design, C_f is the compliance of the de-homogenized design, $\frac{C_f \cdot V_f}{C_{ref} \cdot V_{ref}}$ is the ratio between the performance of the de-homogenized design and the reference solution (lower is better), and t_f is the computation time for the de-homogenization in seconds.

h_c	ε_i	μ_{min}	V_{ref}	C_{ref}	h_f	ε_f	V_f	C_f	$\frac{C_f \cdot V_f}{C_{ref} \cdot V_{ref}}$	t_f [s]
1/30	$20h_i$	0.05	0.2535	106.21	$1/24h_c$	$60h_f$	0.2695	140.16	1.4030	1.28
1/30	$20h_i$	0.05	0.4024	68.58	$1/24h_c$	$60h_f$	0.4329	75.25	1.1804	1.27
1/30	$20h_i$	0.10	0.2568	113.61	$1/24h_c$	$60h_f$	0.2661	167.32	1.5260	1.26
1/30	$20h_i$	0.10	0.4080	69.00	$1/24h_c$	$60h_f$	0.4393	73.68	1.1499	1.33
1/30	$20h_i$	0.20	0.2614	122.86	$1/24h_c$	$60h_f$	0.2665	141.57	1.1747	1.28
1/30	$20h_i$	0.20	0.4165	73.37	$1/24h_c$	$60h_f$	0.4389	77.07	1.1070	1.28
1/30	$10h_i$	0.05	0.2535	106.21	$1/40h_c$	$50h_f$	0.2581	143.65	1.3770	2.23
1/30	$10h_i$	0.05	0.4024	68.58	$1/40h_c$	$50h_f$	0.4201	78.79	1.1993	2.18
1/30	$10h_i$	0.10	0.2568	113.61	$1/40h_c$	$50h_f$	0.2569	149.53	1.3167	2.10
1/30	$10h_i$	0.10	0.4080	69.00	$1/40h_c$	$50h_f$	0.4224	75.16	1.1279	2.19
1/30	$10h_i$	0.20	0.2614	122.86	$1/40h_c$	$50h_f$	0.2566	152.86	1.2214	2.21
1/30	$10h_i$	0.20	0.4165	73.37	$1/40h_c$	$50h_f$	0.4323	77.39	1.0949	2.16

4.1. Network training for specified periodicity

Two different neural networks have been used to generate the results presented in the following sections; one trained for a periodicity equal to a wave-length of $\varepsilon_i = 10h_i$ and one for a wave-length of $\varepsilon_i = 20h_i$. Here h_i refers to the element size on the intermediate mesh \mathcal{T}_i . Both networks have been trained using the two-step training procedure specified in Algorithm 1. In the first step weight factors $\lambda_\omega = 1$, $\lambda_\tau = 1$ and $\lambda_b = 0$ were used, while weight factors $\lambda_\omega = 0$, $\lambda_\tau = 1$ and $\lambda_b = 2$ were used during the second step of training. In all cases the target total variation was set to $\tau = (2/\varepsilon_i)^2$, and the fixed size of the frequency band was set to $b = 4$ pixels. For both the first and second step of training the Adam optimizer with an initial learning rate of $\eta_0 = 2 \cdot e^{-4}$, and $\beta_1 = 0.9$, $\beta_2 = 0.99$ was used. The network was trained for 10 epochs in the first step and 5 epochs in the second step.

A dataset consisting of 20.000 orientation fields sampled as described in Section 3 and encoded using 24 Gaussian channels was used for training. The complete training process, i.e. application of Algorithm 1, for the full 15 epochs using this dataset took ≈ 1 hour on a NVIDIA Titan X GPU, and once trained a forward pass of the network takes only a few milli-seconds on a GPU. To demonstrate that the proposed method also can be used on a standard CPU all computation times presented in the following tables have been performed on a single Intel Core i7 2.6 GHz CPU. If a GPU is available significant speed-ups compared to the stated times can be expected, as the two most expensive operations of the post-processing scheme, the skeletonize and distance transform, can also be performed on a GPU. Wagner [43] reports a 50 times speed-up when performing skeletonization on the GPU, while Zampirolli et al. [44] report a factor five speed-up when performing Euclidean distance transform on the GPU.

4.2. Michell cantilever beam

The extensively studied Michell cantilever beam [45] serves as good starting point for investigating the performance of the neural network-based de-homogenization approach for various parameter configurations. Table 2 shows the results for a fixed homogenization-based resolution of 60×30 with varying periodicity, minimum relative thickness and volume fractions. The performance is measured in terms of compliance and compliance times volume fraction (i.e. specific compliance) normalized by the corresponding reference value from the homogenization-based solution. Note that for small volume fractions, compliance and volume are close to inversely proportional. Consequently, we can use their product (specific compliance) as an approximate means to compare structures with slightly varying volume fractions.

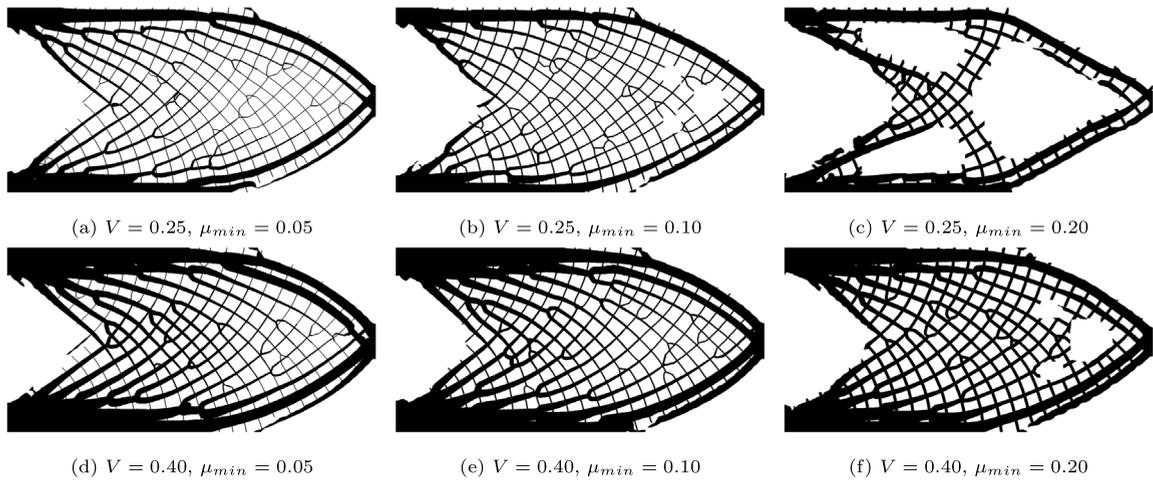


Fig. 10. De-homogenization of a 60×30 Michell cantilever input to a fine mesh of 1440×720 elements with a wave-length of $\varepsilon_f = 60h_f$.

Despite the proposed method not relying on any physical (finite element) modeling in either training or post-processing, the de-homogenized designs are found to perform very well when compared to the reference homogenization-based result. Moreover, the discrepancy is consistently reduced as the wave-length is decreased. That is, on average the de-homogenized solutions performs around 25.7% worse in terms of $C_f \cdot V_f$ for $\varepsilon_i = 20h_i$, and 22.3% worse for $\varepsilon_i = 10h_i$. It is also observed that the proposed method performs best for high volume fractions, which in part can be explained by the fact that structural member thicknesses are better resolved as the volume increases but also because the stiffness of the resulting sub-optimal curved forks is much higher for higher local volume fractions.

Quantitatively, the periodicity does not seem to affect the performance significantly, and thus one could argue that using a larger wave-length should be preferred, as the minimum relative thickness can be resolved on a coarser mesh in this case. However, visual inspection of the designs indicates that the lower wave-length designs seem more robust, for example the diagonal bars in Fig. 10(c) are only resolved by one period in the high wave-length case, whereas multiple periods are present in the low wave-length case, c.f. Fig. 11(c). Boundary artifacts in the form of non-load carrying material is present in many of the designs, especially for larger wave-lengths. One idea to get rid of these could be to use the scheme proposed by Groen et al. [23], in which finite element analyses are used to iteratively identify non-load carrying material, and set these elements to void. Incorporating this post-processing procedure could help reduce the volume fraction used by the neural network-based de-homogenization approach, as it generally tends to overshoot the reference volume fraction slightly, i.e. 5.7% on average for the large wave-length case, and 2.4% for the small wave-length case.

Next, a higher resolution homogenization-based topology optimization result is used as input to the neural network. The results for a 120×60 input de-homogenized onto a 2880×1440 mesh are shown in Appendix A, while the results for a 240×120 input de-homogenized onto a 5760×2880 mesh are shown in Table 3. Here one would expect a higher performance than in the low-resolution input case as the homogenization-based solution itself is of higher quality and individual sub-optimal branches should have less impact on final de-homogenized design. Compared to the low-resolution input case a significant improvement in performance can be observed with the average performance in terms of $C_f \cdot V_f$ being within 13% of the reference value. However, this comes at increased CPU cost as discussed next. Again we note that the obtained objective values can be improved by a post-processing operation to get closer to the values reported in Ref. [23].

4.2.1. Computational cost

The computational costs associated with the method are much lower than the standard TO approach, and outperforms current state-of-the-art de-homogenization methods by several factors. The de-homogenization of the 60×30 homogenization-based solution onto 1440×720 and 2400×1200 meshes took ≈ 1.3 s and ≈ 2.2 s,

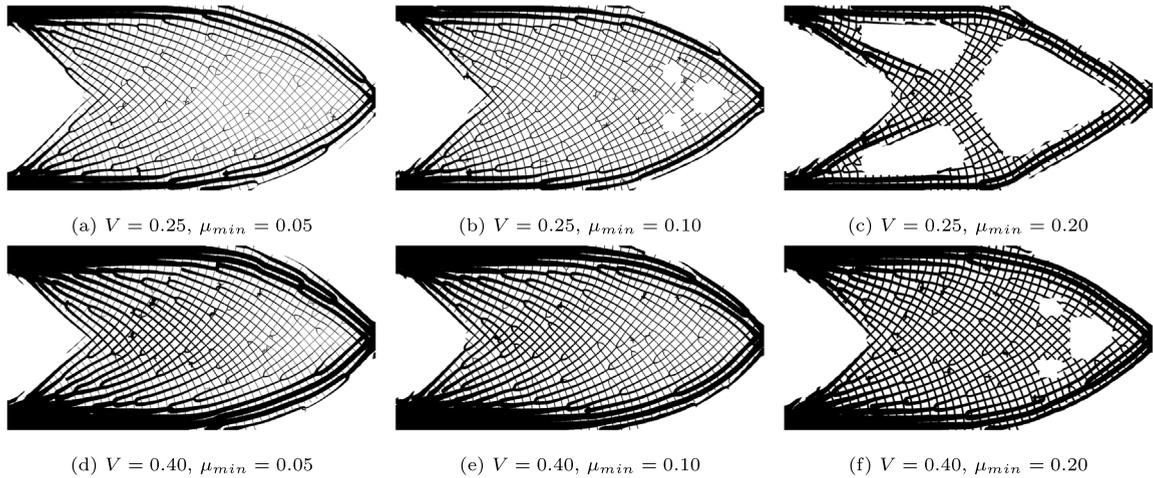


Fig. 11. De-homogenization of a 60×30 Michell cantilever input to a fine mesh of 2400×1200 elements with a wave-length of $\varepsilon_f = 50h_f$.

Table 3

Performance and computational cost of neural network based de-homogenization approach on the Michell cantilever beam for a high-resolution input.

h_c	ε_i	μ_{min}	V_{ref}	\mathcal{C}_{ref}	h_f	ε_f	V_f	\mathcal{C}_f	$\frac{\mathcal{C}_f \cdot V_f}{\mathcal{C}_{ref} \cdot V_{ref}}$	t_f [s]
1/120	$20h_i$	0.05	0.2532	105.98	$1/24h_c$	$60h_f$	0.2587	130.40	1.2575	23.30
1/120	$20h_i$	0.05	0.4023	68.61	$1/24h_c$	$60h_f$	0.4216	73.97	1.1297	23.47
1/120	$20h_i$	0.10	0.2566	111.55	$1/24h_c$	$60h_f$	0.2586	128.85	1.1642	23.41
1/120	$20h_i$	0.10	0.4078	68.94	$1/24h_c$	$60h_f$	0.4237	72.25	1.0885	23.71
1/120	$20h_i$	0.20	0.2572	113.19	$1/24h_c$	$60h_f$	0.2607	120.00	1.0745	24.02
1/120	$20h_i$	0.20	0.4154	71.86	$1/24h_c$	$60h_f$	0.4323	73.72	1.0677	23.55

respectively. At the same resolutions the top88 code¹ took roughly 16s and 87s seconds per iteration respectively. To obtain a design of similar discretization and quality, a continuation strategy would be needed, which typically requires around 1000 iterations to converge [45] leading to a total computation time of 16.000 and 87.000 s respectively. Thus, the neural network-based de-homogenization approach is four orders of magnitudes faster.

In comparison to other de-homogenization approaches; Groen et al. [23] report a computation time of roughly 10 s when solving the least-squares problem on an intermediate mesh to project a 80×40 homogenization-based solution to a 1600×800 mesh. Due to other hardware platform, the computation times are not directly comparable, but both set of computations have been carried out on a modern day laptop, and thus we estimate a decrease in computational cost of around a factor 5 to 10.

A nice property of the proposed method is that the computational cost of all major operations (convolution, bi-linear interpolation, skeletonization and euclidean distance transform) scale linearly in time, making the method very well-suited for generating very high-resolution designs. As an example, for the 5760×2880 mesh (16.5 million elements) direct solvers either become extremely memory consuming or break down entirely, and parallel computing [46] in combination with e.g. multi-grid [47] preconditioners would have to be leveraged in order to perform standard TO. With the proposed method the optimization could be carried out on a coarse scale mesh, e.g. 240×120 elements, using the homogenization method, and afterwards projected onto the high-resolution mesh. All of this could be accomplished in a couple of minutes on a modern day laptop without having to leverage complicated multi-processing frameworks or multi-scale solvers. Fig. 13 shows the average computation time for the six designs in Fig. 12, along with a breakdown of the computation time for different sub-components of the framework. Here the loss function evaluation, solidify branches, network forward pass, skeletonize and distance transform are the most costly components. The solidify branches component is mostly a safety measure, to make

¹ <https://www.topopt.mek.dtu.dk/Apps-and-software/Efficient-topology-optimization-in-MATLAB>

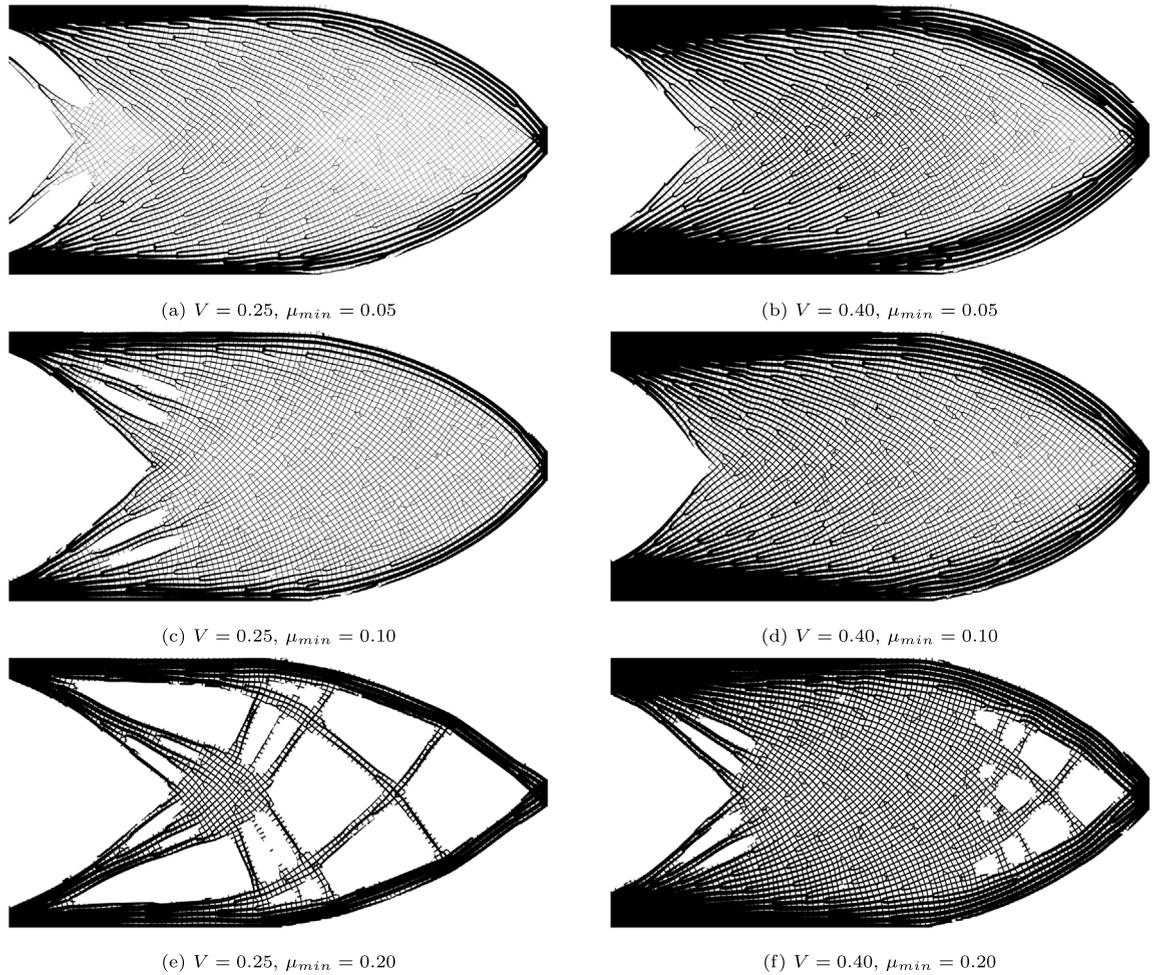


Fig. 12. De-homogenization of a 240×120 Michell cantilever input to a fine mesh of 5760×2880 elements with a wave-length of $\varepsilon_f = 60h_f$.

sure that branches are always solid, but is often not needed as the two-step training procedure is enough to assure this in most cases. As this is the only component in the post-processing procedure relying on a loss function evaluation dropping this component could save up to 40% of the computational cost. The three remaining costly components can, as mentioned earlier, all be run on a GPU for a significant speed-up. In terms of memory consumption the proposed method is able to handle inputs up to a resolution of 800×800 on a single 11GB GPU. This corresponds to a de-homogenized design with a resolution of 32000×32000 when using $\varepsilon_i = 10$ which should be more than enough for representing any currently manufacturable design in 2D. For 3D purposes the framework can with no further theoretical development be extended to multi-GPU usage by wrapping the model with a distribution module, such as `tf.distribute`² in Tensorflow or `torch.nn.DataParallel`³ in PyTorch.

Considering these significant advantages in terms of computational cost a loss in performance (before further post-processing) of 7–25% is not immense, and in the low-resolution case where the de-homogenization only takes a few seconds the proposed approach could serve as a real-time evaluation tool for a discrete and manufacturable version of the homogenized-based design. In the high-resolution case the de-homogenized design could serve as a good starting guess for high-resolution shape or topology optimization, thus saving a large amount of iterations.

² https://www.tensorflow.org/api_docs/python/tf/distribute

³ <https://pytorch.org/docs/stable/generated/torch.nn.DataParallel.html>

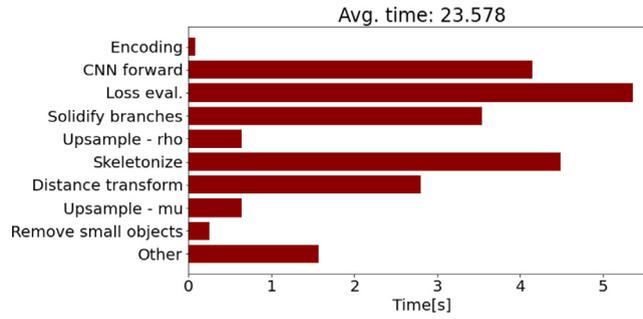


Fig. 13. Average computation time for different sub-components of the de-homogenization framework when run on six high-resolution (5760×2880) examples.

Table 4

Performance and computational cost of the neural network based de-homogenization approach on the double-clamped beam example.

h_c	ε_i	μ_{min}	V_{ref}	C_{ref}	h_f	ε_f	V_f	C_f	$\frac{C_f \cdot V_f}{C_{ref} \cdot V_{ref}}$	t_f [s]
1/50	$20h_i$	0.05	0.2510	25.89	$1/24h_c$	$60h_f$	0.2590	31.35	1.2492	7.69
1/50	$20h_i$	0.10	0.2538	27.11	$1/24h_c$	$60h_f$	0.2619	32.51	1.2369	7.12
1/50	$10h_i$	0.05	0.2510	25.89	$1/40h_c$	$50h_f$	0.2568	31.50	1.2446	15.80
1/50	$10h_i$	0.10	0.2538	27.11	$1/40h_c$	$50h_f$	0.2574	31.12	1.1657	14.70

Table 5

Performance and computational cost of neural network based de-homogenization approach on the L-shaped beam.

h_c	ε_i	μ_{min}	V_{ref}	C_{ref}	h_f	ε_f	V_f	C_f	$\frac{C_f \cdot V_f}{C_{ref} \cdot V_{ref}}$	t_f [s]
1/80	$20h_i$	0.10	0.2517	94.15	$1/24h_c$	$60h_f$	0.2758	96.29	1.1207	4.96
1/80	$10h_i$	0.10	0.2517	94.15	$1/40h_c$	$50h_f$	0.2742	96.52	1.1168	7.53

4.3. Double-clamped beam

As a second example, the proposed method is applied to the double-clamped beam with center loading. An input of size 200×50 is considered and de-homogenized at two different wave-lengths with the fine-scale mesh size chosen according to the minimum relative thickness. The results are shown in Table 4, and the corresponding designs in Fig. 14. Here the average performance is 22.4% worse than the reference case, and the volume constraint is violated by an average of 2.3%. The best performance is achieved for the small wave-length case, but at a cost of double the computation time.

4.4. L-shaped beam

As a last example the L-shaped beam is considered. This example presents a more challenging case due to the singularity at the sharp corner in the design domain. Directly solving the optimization problem in Eq. (1) leads to small regions near the sharp corner, in which the angles are turned 90° compared to neighboring elements. One way to remedy this issue is the vector field combing method presented in Stutz et al. [42] which provides a consistent labeling of the lamination orientations. Table 5 and Fig. 15 show the results for the L-shaped beam using the combed input fields. Here it can be seen that higher periodicity only leads to a slightly better performance with both designs being within roughly 12% of the reference solution.

5. Conclusions and future work

An unsupervised deep learning-based de-homogenization framework capable of generating high-resolution, manufacturable, and continuous microstructures has been presented. Compared to previous deep learning-based

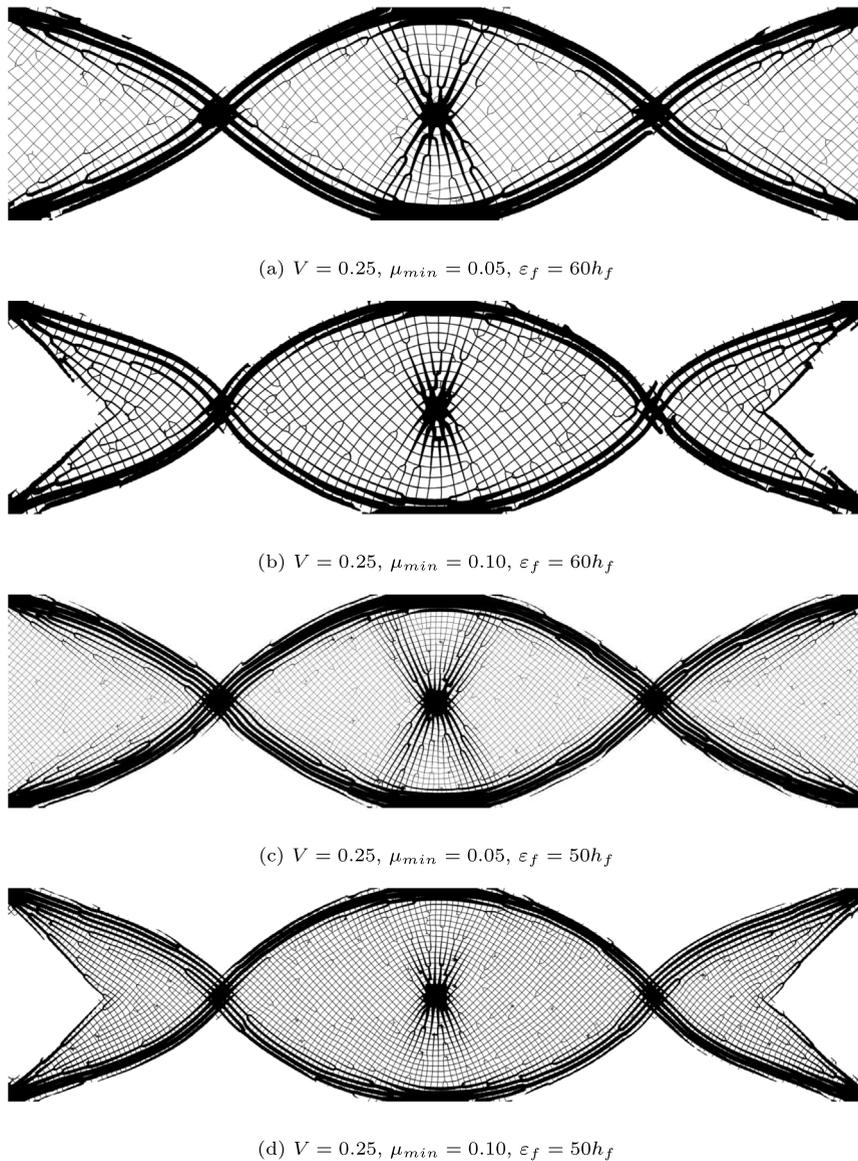


Fig. 14. De-homogenization of a 200×50 double-clamped beam input to a fine mesh of 4800×1200 (a and b) and 8000×2000 (c and d).

topology optimization approaches the proposed method does not seek to perform topology optimization in an end-to-end manner, but instead utilizes the solution from a traditional homogenization-based topology optimization as input. This leads to the, perhaps, most important feature of the proposed methodology: it is, to a large extent, insensitive to domain size and shape, loading, and boundary conditions. The proposed fully convolutional neural network approach is capable of synthesizing continuous and periodic microstructures, while the input, in the form of an orientation field, guides the network towards a mechanically sound design. A subsequent post-processing scheme assigns an appropriate thickness to each member of the final design, and ensures manufacturability by imposing a minimum relative thickness. Training of the neural network is extremely inexpensive as synthetic orientation fields can be used as training data, and generalizes to a wide range of problems, as the underlying load and boundary conditions are implicitly enforced by the homogenization-based input. Numerical examples show that the proposed

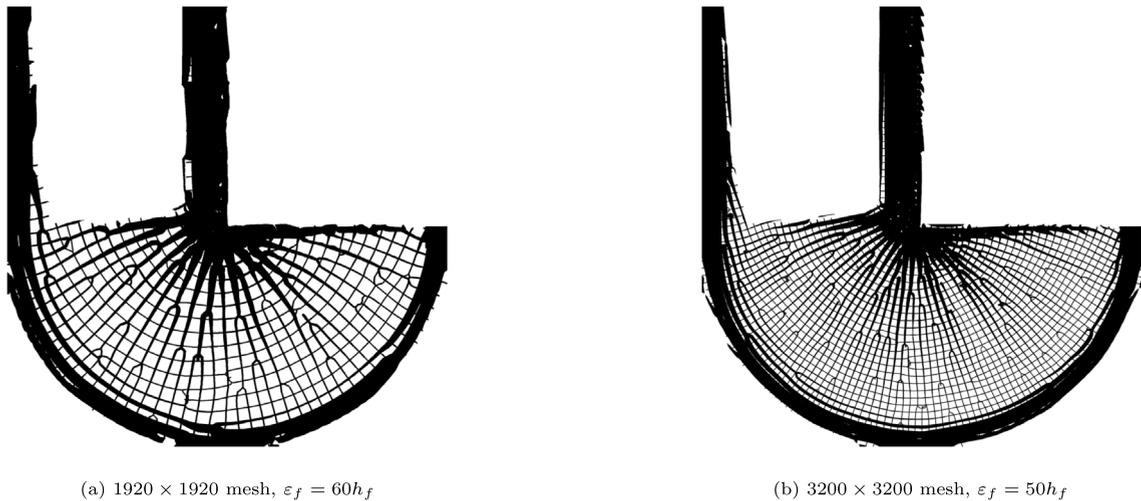


Fig. 15. De-homogenization of the L-shaped beam with a combed orientation field as input to the neural network. $V = 0.25$, $\mu_{min} = 0.10$.

method is a factor 5 to 10 faster than current state-of-the-art de-homogenization approaches and has a performance within 7–25% of the reference solution (before further post-processing). While only 2D examples are shown the very low computational cost of the method makes it an obvious candidate for 3D applications. Furthermore, as all operations in the proposed de-homogenization pipeline support GPU computations, a GPU implementation could pave the way for interactive de-homogenization, i.e. projection during the optimization, in 2D cases.

There are several ways upon which the current method could be improved. For example, lamination widths could be used directly as input to the neural network, instead of being applied in the post-processing scheme. This would allow the neural network to trade a high error in inconsequential parts of the domain, i.e. regions with very little material, for a lower error in solid parts of the domain. As of now, the proposed method also tends to slightly overshoot the prescribed volume fraction. This problem could be remedied by iteratively performing a finite element analysis to identify non-loaded carrying members in the final design and remove them. Furthermore, a graph-based post-processing scheme could be used to identify U-shaped branches and turn them into V-shaped branches which are known to be a mechanically better design when designing for minimum compliance. Finally, the iso-contour of the design could be extracted directly from the distance transform utilized in the post-processing scheme, such that the design can be described using a conforming mesh. This would yield a smoother boundary representation while allowing a more efficient finite element analysis during post-processing and subsequent shape optimization. An obvious idea, to be studied in future work, is the extension to 3D, which is expected to result in even larger time savings than reported here for the 2D case.

Apart from the immediate extensions and improvements for single load compliance minimization discussed above, we note that the proposed approach is general and applies to all kinds of other problems and formulations that require mapping of directional fields to periodic fine scale structures. A first extension could be the mapping of stress constrained problems which were solved using rank-2 based homogenization in the original work by Duysinx and Bendsøe [48]. The results from their approach can directly be mapped to fine scale structures using the suggested approach.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors would like to acknowledge the support of the Villum Foundation, Denmark through the Villum Investigator Project InnoTop. Gratitude should also be expressed to Anders Nymark Christensen and Morten

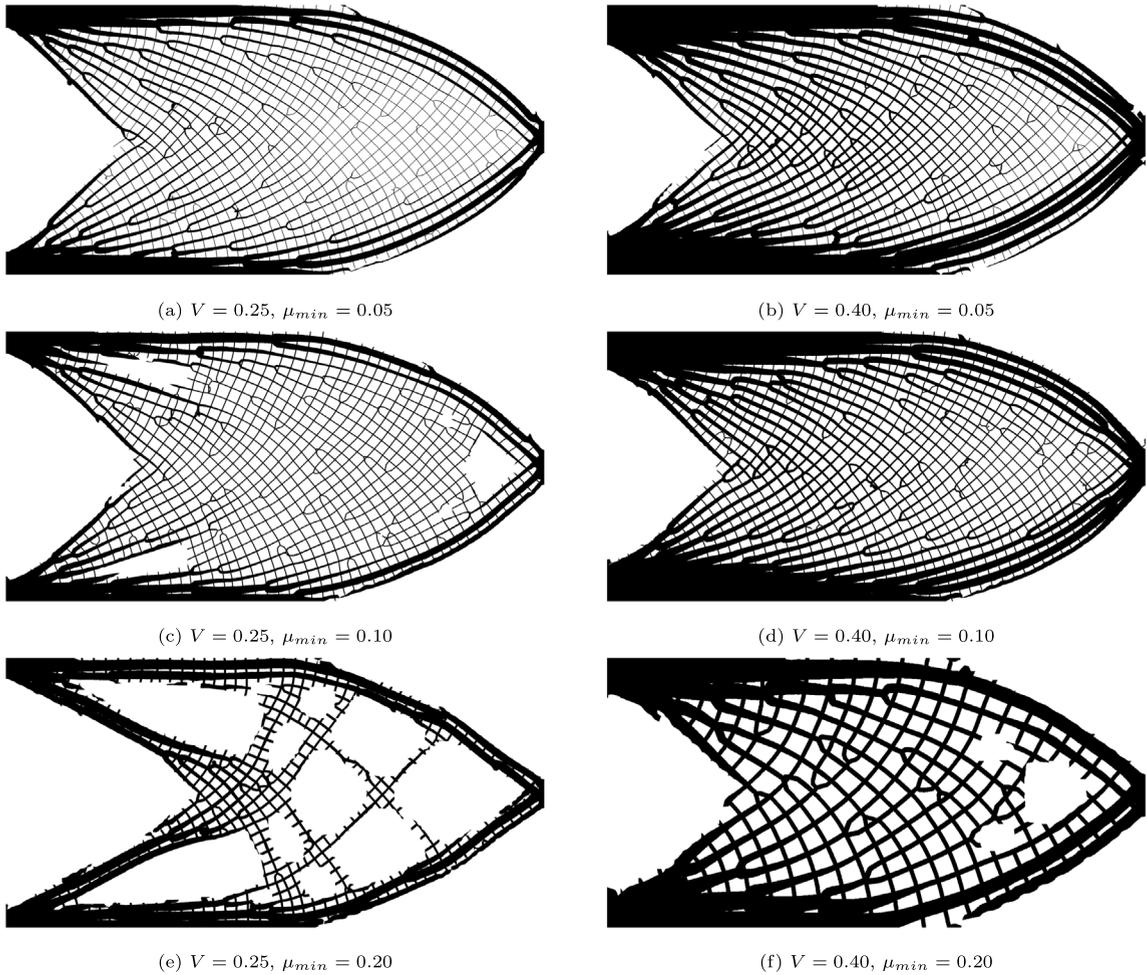


Fig. A.1. De-homogenization of a 120×60 Michell cantilever input to a fine mesh of 2880×1440 elements with a wave-length of $\varepsilon_f = 60h_f$.

Hannemose for valuable discussions on the deep learning parts of the project. Finally, the authors would like to thank Jeroen Peter Groen and Reinier Giele for providing the homogenization-based topology optimization code and to Florian Stutz for providing the code for angular field combing.

Appendix A. Michell cantilever - 120×60 designs

See [Table A.1](#) and [Fig. A.1](#).

Appendix B. Network architecture

The neural network consists of 9 convolutional layers, four ResNet blocks and three upsampling layers. The two first layers use a kernel size of 7 and 5 respectively, while all other layers use a kernel size of 3. Between each convolutional layer batch normalization is applied, and ReLU is used the activation function in all layers except the last, which uses a Sigmoid. For upsampling nearest neighbors is used, as this is less prone to checkerboard artifacts than learned upsampling filters. The network architecture is summarized in [Table B.1](#).

Table A.1

Performance and computational cost of neural network based de-homogenization approach on the Michell cantilever beam for a 120×60 input.

h_c	ε_i	μ_{min}	V_{ref}	C_{ref}	h_f	ε_f	V_f	C_f	$\frac{C_f \cdot V_f}{C_{ref} \cdot V_{ref}}$	t_f [s]
1/60	$20h_i$	0.05	0.2534	105.92	$1/24h_c$	$60h_f$	0.2608	128.52	1.2489	5.54
1/60	$20h_i$	0.05	0.4024	68.58	$1/24h_c$	$60h_f$	0.4219	74.36	1.1366	5.52
1/60	$20h_i$	0.10	0.2567	112.42	$1/24h_c$	$60h_f$	0.2599	134.91	1.2148	5.85
1/60	$20h_i$	0.10	0.4080	68.95	$1/24h_c$	$60h_f$	0.4309	71.89	1.1010	5.38
1/60	$20h_i$	0.20	0.2593	117.60	$1/24h_c$	$60h_f$	0.2601	175.84	1.4998	5.61
1/60	$20h_i$	0.20	0.4163	72.89	$1/24h_c$	$60h_f$	0.4391	74.60	1.0798	5.45

Table B.1

Convolutional neural network architecture.

Layer	Description
Input	\mathbb{R}^N $H \times W$
1	7×7 conv, $C = 32$, BatchNorm, ReLU
2	5×5 conv, $C = 64$, BatchNorm, ReLU
3	$(3 \times 3$ conv, $C = 64$, ResNet block) $\times 4$
4	Nearest neighbor upsampling, scale=2
5	3×3 conv, $C = 64$, BatchNorm, ReLU
6	3×3 conv, $C = 64$, BatchNorm, ReLU
7	Nearest neighbor upsampling, scale=2
8	3×3 conv, $C = 64$, BatchNorm, ReLU
9	3×3 conv, $C = 32$, BatchNorm, ReLU
10	Nearest neighbor upsampling, scale=2
11	3×3 conv, $C = 32$, BatchNorm, ReLU
12	3×3 conv, $C = 32$, BatchNorm, ReLU
13	3×3 conv, $C = 1$, Sigmoid

References

- [1] M.P. Bendsøe, O. Sigmund, *Topology Optimization: Theory, Methods and Applications*, second ed., Springer, 2004.
- [2] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (7553) (2015) 436–444, <http://dx.doi.org/10.1038/nature14539>.
- [3] S. Hoyer, J. Sohl-Dickstein, S. Greydanus, Neural reparameterization improves structural optimization, Pre-print <https://arxiv.org/abs/1909.04240>.
- [4] D. Wang, C. Xiang, Y. Pan, A. Chen, X. Zhou, Y. Zhang, A deep convolutional neural network for topology optimization with perceptible generalization ability, *Eng. Optim.* (2021) 1–16, <http://dx.doi.org/10.1080/0305215X.2021.1902998>.
- [5] R. Cang, H. Yao, Y. Ren, One-shot generation of near-optimal topology through theory-driven machine learning, *Comput. Aided Des.* 109 (2019) 12–21, <http://dx.doi.org/10.1016/j.cad.2018.12.008>.
- [6] Y. Yu, T. Hur, J. Jung, I.G. Jang, Deep learning for determining a near-optimal topological design without any iteration, *Struct. Multidiscip. Optim.* 59 (3) (2019) 787–799, <http://dx.doi.org/10.1007/s00158-018-2101-5>.
- [7] A. Chandrasekhar, K. Suresh, Tounn: Topology optimization using neural networks, *Struct. Multidiscip. Optim.* 63 (3) (2021) 1135–1149, <http://dx.doi.org/10.1007/s00158-020-02748-4>.
- [8] Z. Nie, T. Lin, H. Jiang, L.B. Kara, Topologygan: Topology optimization using generative adversarial networks based on physical fields over the initial domain, *J. Mech. Des.* 143 (3) (2021) <http://dx.doi.org/10.1115/1.4049533>.
- [9] N.A. Kallioras, G. Kazakis, N.D. Lagaros, Accelerated topology optimization by means of deep learning, *Struct. Multidiscip. Optim.* 62 (3) (2020) 1185–1212, <http://dx.doi.org/10.1007/s00158-020-02545-z>.
- [10] Q. Lin, J. Hong, Z. Liu, B. Li, J. Wang, Investigation into the topology optimization for conductive heat transfer based on deep learning approach, *Int. Commun. Heat Mass Transfer* 97 (2018) 103–109, <http://dx.doi.org/10.1016/j.icheatmasstransfer.2018.07.001>.
- [11] B. Li, C. Huang, X. Li, S. Zheng, J. Hong, Non-iterative structural topology optimization using deep learning, *Comput. Aided Des.* 115 (2019) 172–180, <http://dx.doi.org/10.1016/j.cad.2019.05.038>.
- [12] B. Li, W. Tang, S. Ding, J. Hong, A generative design method for structural topology optimization via transformable triangular mesh algorithm, *Struct. Multidiscip. Optim.* 62 (3) (2020) 1159–1183, <http://dx.doi.org/10.1007/s00158-020-02544-0>.
- [13] D.W. Abueidda, M. Almasri, R. Ammourah, U. Ravaioli, I.M. Jasiuk, N.A. Sobh, Prediction and optimization of mechanical properties of composites using convolutional neural networks, *Compos. Struct.* 227 (2019) 111264, <http://dx.doi.org/10.1016/j.compstruct.2019.111264>.
- [14] R.K. Tan, N.L. Zhang, W. Ye, A deep learning-based method for the design of microstructural materials, *Struct. Multidiscip. Optim.* 61 (4) (2020) 1417–1438, <http://dx.doi.org/10.1007/s00158-019-02424-2>.
- [15] L. Wang, Y.-C. Chan, F. Ahmed, Z. Liu, P. Zhu, W. Chen, Deep generative modeling for mechanistic-based learning and design of metamaterial systems, *Comput. Methods Appl. Mech. Engrg.* 372 (2020) 113377, <http://dx.doi.org/10.1016/j.cma.2020.113377>.

- [16] H.T. Kollmann, D.W. Abueidda, S. Koric, E. Guleryuz, N.A. Sobh, Deep learning for topology optimization of 2d metamaterials, *Mater. Des.* 196 (2020) 109098, <http://dx.doi.org/10.1016/j.matdes.2020.109098>.
- [17] C.T. Chen, G.X. Gu, Generative deep neural networks for inverse materials design using backpropagation and active learning, *Adv. Sci.* 7 (5) (2020) 1902607, <http://dx.doi.org/10.1002/adv.201902607>.
- [18] S. Oh, Y. Jung, S. Kim, I. Lee, N. Kang, Deep generative design: Integration of topology optimization and generative models, *Trans. ASME, J. Mech. Des.* 141 (11) (2019) e4044229, <http://dx.doi.org/10.1115/1.4044229>.
- [19] D. Ulyanov, A. Vedaldi, V. Lempitsky, Deep image prior, *Int. J. Comput. Vis.* 128 (2020) 1867–1888, <http://dx.doi.org/10.1007/s11263-020-01303-4>.
- [20] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, W. Shi, Photo-realistic single image super-resolution using a generative adversarial network, in: *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017, pp. 105–114, <http://dx.doi.org/10.1109/CVPR.2017.19>.
- [21] H. Chi, Y. Zhang, T.L.E. Tang, L. Mirabella, L. Dalloro, L. Song, G.H. Paulino, Universal machine learning for topology optimization, *Comput. Methods Appl. Mech. Engrg.* 375 (2021) 112739, <http://dx.doi.org/10.1016/j.cma.2019.112739>.
- [22] O. Pantz, K. Trabelsi, A post-treatment of the homogenization method for shape optimization, *SIAM J. Control Optim.* 47 (3) (2008) 1380–1398, <http://dx.doi.org/10.1137/070688900>.
- [23] J.P. Groen, O. Sigmund, Homogenization-based topology optimization for high-resolution manufacturable micro-structures, *Internat. J. Numer. Methods Engrg.* 113 (8) (2018) 1148–1163, <http://dx.doi.org/10.1002/nme.5575>.
- [24] G. Allaire, P. Geoffroy-Donders, O. Pantz, Topology optimization of modulated and oriented periodic microstructures by the homogenization method, *Comput. Math. Appl.* 78 (7) (2019) 2197–2229, <http://dx.doi.org/10.1016/j.camwa.2018.08.007>.
- [25] J.P. Groen, F.C. Stutz, N. Aage, J.A. Barentzen, O. Sigmund, De-homogenization of optimal multi-scale 3d topologies, *Comput. Methods Appl. Mech. Engrg.* 364 (2020) 112979, <http://dx.doi.org/10.1016/j.cma.2020.112979>.
- [26] M.P. Bendsøe, Generating optimal topologies in structural design using a homogenization method, *Comput. Methods Appl. Mech. Eng.*
- [27] E. Shelhamer, J. Long, T. Darrell, Fully convolutional networks for semantic segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (4) (2017) 7478072, <http://dx.doi.org/10.1109/TPAMI.2016.2572683>, 640–651.
- [28] M. Nobel-Jørgensen, N. Aage, A. Christiansen, T. Igarashi, J. Barentzen, O. Sigmund, 3d interactive topology optimization on hand-held devices, *Struct. Multidiscip. Optim.* 51 (6) (2015) 1385–1391, <http://dx.doi.org/10.1007/s00158-014-1214-8>.
- [29] M.P. Bendsøe, Optimal shape design as a material distribution problem, *J. Struct. Optim.* 1 (1989) 193–202.
- [30] M.P. Bendsøe, O. Sigmund, Material interpolation schemes in topology optimization, *Arch. Appl. Mech.* 69 (9–10) (1999) 635–654.
- [31] M. Baandrup, O. Sigmund, H. Polk, N. Aage, Closing the gap towards super-long suspension bridges using computational morphogenesis, *Nature Commun.* 11 (2020) 2735.
- [32] R. Giele, J. Groen, N. Aage, C. Andreassen, O. Sigmund, On approaches for avoiding low stiffness regions in variable thickness sheet and homogenization-based topology optimization, *Struct. Multidiscip. Optim.* 64 (2021) 39–52.
- [33] F. Wang, B.S. Lazarov, O. Sigmund, On projection methods, convergence and robust formulations in topology optimization, *Struct. Multidiscip. Optim.* 43 (2011) 767–784, <http://dx.doi.org/10.1007/s00158-010-0602-y>.
- [34] P. Norvig, S.J. Russell, *Artificial Intelligence: A Modern Approach*, Pearson, 2010.
- [35] Z. Ghahramani, Unsupervised learning, in: *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 3176, 2004, pp. 72–112, http://dx.doi.org/10.1007/978-3-540-28650-9_5.
- [36] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition 2016-*, 2016, 7780459, <http://dx.doi.org/10.1109/CVPR.2016.90>.
- [37] A. Mordvintsev, C. Olah, M. Tyka, Inceptionism: Going deeper into neural networks, 2015, <https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>.
- [38] J. Groen, J. Wu, O. Sigmund, Homogenization-based stiffness optimization and projection of 2d coated structures with orthotropic infill, *Comput. Methods Appl. Mech. Engrg.* 349 (2019) 722–742, (also [arXiv:1808.04740](https://arxiv.org/abs/1808.04740) August 2018).
- [39] F.J. Harris, On the use of windows for harmonic analysis with the discrete fourier transform, *Proc. IEEE* 66 (1) (1978) 51–83, <http://dx.doi.org/10.1109/PROC.1978.10837>.
- [40] E. Träff, O. Sigmund, J. Groen, Simple single-scale microstructures based on rank-3 optimal laminates, *Struct. Multidiscip. Optim.* 59 (4) (2019) 1021–1031.
- [41] T.Y. Zhang, C.Y. Suen, A fast parallel algorithm for thinning digital patterns, *Commun. ACM* 27 (3) (1984) 236–239, <http://dx.doi.org/10.1145/357994.358023>.
- [42] F.C. Stutz, J.P. Groen, O. Sigmund, J.A. Barentzen, Singularity aware de-homogenization for high-resolution topology optimized structures, *Struct. Multidiscip. Optim.* 62 (5) (2020) 2279–2295, <http://dx.doi.org/10.1007/s00158-020-02681-6>.
- [43] M.G. Wagner, Real-time thinning algorithms for 2d and 3d images using gpu processors, *J. Real-Time Image Process.* 17 (5) (2020) 1255–1266, <http://dx.doi.org/10.1007/s11554-019-00886-7>.
- [44] F.D.A. Zampiroli, L. Filipe, A fast cuda-based implementation for the euclidean distance transform, in: *Proceedings - 2017 International Conference on High Performance Computing and Simulation, HPCS 2017*, <http://dx.doi.org/10.1109/HPCS.2017.123>.
- [45] O. Sigmund, N. Aage, E. Andreassen, On the (non-)optimality of michell structures, *Struct. Multidiscip. Optim.* 54 (2) (2016) 361–373, <http://dx.doi.org/10.1007/s00158-016-1420-7>.
- [46] N. Aage, E. Andreassen, B.S. Lazarov, Topology optimization using petsc: An easy-to-use, fully parallel, open source topology optimization framework, *Struct. Multidiscip. Optim.* 51 (3) (2015) 565–572, <http://dx.doi.org/10.1007/s00158-014-1157-0>.
- [47] O. Amir, N. Aage, B.S. Lazarov, On multigrid-cg for efficient topology optimization, *Struct. Multidiscip. Optim.* 49 (5) (2014) 815–829, <http://dx.doi.org/10.1007/s00158-013-1015-5>.
- [48] P. Duysinx, M.P. Bendsøe, Topology optimization of continuum structures with local stress constraints, *Internat. J. Numer. Methods Engrg.* 43 (8) (1998) 1453–1478.